



Naval Research Advisory Committee Report

Future Naval Use of COTS Networking Infrastructure

July 2009

Report Documentation Page			<i>Form Approved OMB No. 0704-0188</i>	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE JUL 2009	2. REPORT TYPE	3. DATES COVERED 00-00-2009 to 00-00-2009		
4. TITLE AND SUBTITLE Future Naval Use of COTS Networking Infrastructure		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Advisory Committee, 875 North Randolph Street Suite 1230, Arlington, VA, 22203-1995		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 68
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified		

This report is a product of the U.S. Naval Research Advisory Committee (NRAC) Panel on Future Naval Use of COTS Networking Infrastructure, opinions, recommendations, and/or conclusions contained in this report are those of the NRAC Panel and do not necessarily represent the official position of the U.S. Navy, or the Department of Defense.

Table of Contents

Executive Summary	Pages 5-6
Main Body of Report	Pages 7-53
Terms of Reference	
Fact Finding	
Naval IT Programs	
Why this Matters	
Network Evolution	
Resource Shaping	
Cloud Computing	
Impact	
Findings	
Action Items	
Take-Aways	
Appendices:	Pages 54-68
A. Terms of Reference Document	
B. Sample Definitions of Cloud Computing	
C. Panel Members	
D. Fact-Finding Contributors	
E. Acronyms	

This page intentionally left blank

Executive Summary

This study was conducted by the Naval Research Advisory Committee under the sponsorship of the Deputy Chief of Naval Operations for Communication Networks (N6). It explores the future of commercial networking architectures and the ways in which these architectures might be adapted to Naval use, based on the current pace and direction of their development. Today, the Navy and Marine Corps incorporate commercial networking architectures and this will continue into the future. In fact, the impact of developing proprietary networking architectures – on capital, maintenance, and training expenses – would be cost prohibitive. The fundamental question to be answered by this study is, “How can the Navy and Marine Corps leverage emerging commercial networking architectures to improve operational effectiveness while keeping the cost in an affordable range?”

The report is divided broadly into four sections. The first establishes the reasons why the study matters to the Navy. It sets the stage with an example from the Bechtel Corporation, a large company that deploys task units into remote areas where bandwidth is limited and connectivity is intermittent, i.e., conditions that resemble those in many Naval operating areas. The second section addresses the evolution of networking architectures from the original productivity-focused computing systems in the mid-1970s to the latest next-generation networking architectures. It begins with the centralized architectures (circa 1970-1985), moves through the era of the client-server architectures introduced in the mid-1980’s and still the dominant networking architecture today, and shows the current transition to the new architectures that are characterized by virtual execution and central management. It then focuses on the evolution of resource sharing, illustrating the advances of the newer architectures over the current, and introduces the concept and definition of Cloud Computing. While Cloud Computing is developing in many variations – including Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS), among others – the overall concept of shared resources across accessible, virtual, scalable systems is common to all. The second section ends with the impact that Cloud architectures are having on the way digital information and communications are handled today.

The third section provides a transition from the discussion of the technical aspects of Cloud Computing, as it is being implemented in the commercial sector, to the technical gaps between the priorities of the commercial development and Naval requirements. The identified technical gaps were developed by comparing the priorities of the commercial developers to the requirements of the Naval Networking Environment goals and features. Also, operational requirements were derived from meetings with the Commander Third Fleet N6 staff and a review of recent after action reports from deployed strike groups. The gaps are grouped under two headings: 1) security in the Cloud and 2) bandwidth and connectivity in the Cloud. These are clearly the dominant gaps that must be addressed by science and technology before the Navy and Marine Corps can embrace Cloud Computing as their operational standard.

The final section contains the Panel's findings and recommendations. The "take-away" points are what the Panel considers the most important. They are:

- Cloud Computing is here to stay,
- Engage the Cloud community to ensure Naval needs are incorporated into evolving standards,
- Establish Cloud pilot project(s) for non-combat services,
- Focus research and development efforts on:
 - Securing the virtualization layer,
 - Developing data links that enable Cloud architectures,
 - Developing Cloud performance models to simulate network performance in various conditions.

A cautionary note for readers of this report: studies in IT offer a great opportunity to explore the leading edge – and to become outdated in the process. The pace of IT development is well beyond one's ability to maintain complete knowledge of ongoing developments. As an example – since the first draft of this report on July 1, 2009 – the commercial sector has begun to address at least one (and probably many more) of the issues the Panel noted as a gap to Naval applications: the problem of time delays in communication or information links, i.e., latency.



Terms of Reference

- Study the Navy's use of commercial architectures, software, and hardware.
- Examine the related emerging networking approaches under development in the commercial world.
- Examine the development and operational practices associated with the emerging approaches.
- Suggest strategies for leveraging ongoing COTS investment in future Naval networks:
 - In light of dramatically changing Naval bandwidth availability, uncertain connectivity, and large latencies
 - Within a global supply chain
- Recommend S&T investments to adapt the emerging networking approaches to Naval requirements.

This slide shows an overview of the study Terms of Reference (TOR), provided in full form in Appendix A. Our study examines the emerging networking developments in the commercial Information Technology (IT) world, as well as the architectures and operational practices associated with them. And it suggests strategies for leveraging ongoing civilian investment for Navy needs – identifying S&T efforts to adapt those commercial technologies to Navy operational and administrative requirements. One should note that the TOR is very broad.

Specific Tasking:

- Compare and contrast the Navy/Marine Corps needs to maintain secure network functionality in the face of dramatically changing bandwidth availability, uncertain connectivity, and large latencies with capabilities offered by existing and emerging commercial technology;

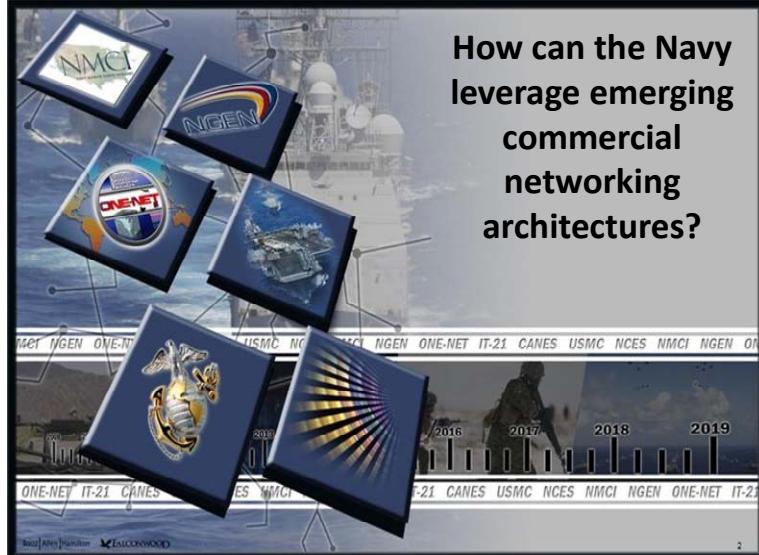
- Explore how the integrity of Navy/Marine Corps networks can be assured with commercially developed components, e.g., when personnel developing commercial software will most certainly include non-US citizens; and,
- Review current Navy S&T and develop a set of actionable recommendations for new investments as well as changes to current investments that must be made by the Navy in its S&T portfolio to exploit promising commercial networking technologies.



Fact Finding



The Panel had discussions and received briefings from a number of organizations that provide or use IT services, represented by their logos as shown above. These include a number that operate “in the Cloud.” The top half of the chart shows providers/users in the commercial world – the lower half shows government and military organizations.



Cloud Computing can serve as a significant enabler of the converged enterprise vision, as set forth in the 2016 Naval Network Environment (NNE) vision and roadmap documents. The NNE is a federation of enterprise services, internet protocol-based networks, and enabling infrastructure bounded by common enterprise architecture, standards, and governance consistent with the network operations (NetOps) concept. The NNE will provide ready, secure access to data and services across the Naval enterprise and interoperability with the networks of other services and agencies.

The 16 objectives of the NNE enabled by Cloud Computing concepts and technologies are:

1. Ability to log on to the network and securely access data from anywhere in the NNE;
2. Elimination of legacy networks and all users on an approved NNE network;
3. Common NNE governance that results in secure, interoperable networks and improved return on investment;
4. Highly secure networks and robust Computer Network Defense (CND);
5. Fully integrated NETOPS capability;

6. Enforcement of target NNE Architecture that leverages enterprise IT initiatives;
7. Common enterprise applications to include a common directory service;
8. Asset management and enterprise licensing;
9. Innovative acquisition strategies and improved financial visibility;
10. Ability to seamlessly collaborate across the NNE and with other services, agencies, and partners;
11. Enterprise content management;
12. Flexible, adaptable NNE capable of responding to continually evolving operations and network threats;
13. Data strategy;
14. Government control of the NNE design and operation of the network;
15. Performance management and continuous process improvement;
16. Development and deployment of a robust and scalable network infrastructure capable of IP-based data.

The Navy/Marine Corps Intranet (NMCI) provides approximately 700,000 Navy and Marine Corps user accounts on 340,000 “seats” servicing over 3,000 locations across the continental United States, Hawaii, Cuba, Guam, Japan, and Puerto Rico. This contractor-owned/contractor-operated network was designed to provide universal access to integrated voice, video, and data communications and a common computing environment, playing a critical role in improving security across the enterprise.

In addition to the OCONUS sites served by NMCI, the OCONUS Navy Enterprise Network (ONE-NET), including the Base Level Information Infrastructure (BLII) efforts, provides roughly 41,000 users at shore installations overseas, a single integrated network with a full range of services, and a centralized control authority. It is a government-owned, government-operated, Navy enterprise network that delivers centralized IT, improved security, standard configurations, and increased service levels to overseas commands.

For the afloat forces, the Information Technology for the 21st Century (IT-21) program provides networking capabilities across the fleet. Examples of these are local area networks such as the Integrated Shipboard Network System (ISNS); communications services such as the Global Command and Control System - Maritime (GCCS-M); routing services such as the Automated Digital Network System (ADNS); communications on Extremely High Frequency (EHF), Super High Frequency (SHF) and Ultra High Frequency (UHF) satellite systems; and the shore-based infrastructure that supports global operations (e.g., Tactical Switching).

For the Marines, the Marine Corps Enterprise Network (MCEN) is a portfolio of IT programs that provide network services to Marines in CONUS, OCONUS, and deployed Marine Air-Ground Task Forces (MAGTFs). MCEN utilizes NMCI and IT-21 capabilities as well as its own internal tactical function.

In addition, the current Naval Networking Environment is also comprised of over 500 legacy networks that support specific functions within DON or cross-service organizations. Because the legacy networks are also usually associated with legacy applications, work continues to reduce and replace legacy applications with new applications that are not dependent on specific hardware and software. Examples of these networks are the DON Research and Development and Test and Evaluation networks, which carry large volumes of telemetry, video, and/or computer-aided design data; the Tri-Care and Service Medical networks; the Training Commands and Service Academy classroom training and education networks; and the networks that support the Service's Recruiting Command functions.

Within the 62 acquisition programs that make up the NNE, the two most affected by Cloud Computing are the Next Generation Enterprise Network (NGEN) and Consolidated Afloat Network Enterprise Services (CANES). These initiatives make up the bulk of the NNE and serve as the two “transformational” opportunities for leveraging the capabilities of Cloud Computing.

Existing Naval IT expenditures outside of the enterprise IT programs are identified and aligned through the Navy's Cyber Asset Reduction and Security (CARS) program. It has been proposed that this program continue in order to provide continuous process improvement in support of NGEN and CANES under the Information Technology Infrastructure Library (ITIL) model. In addition, the Navy should consider the migration of capabilities and requirements to computing services hosted in public “Clouds” for non-sensitive data, i.e., those not critical to the Navy’s warfighting mission. Private, government-owned “Clouds” hosted at the Navy Defense Enterprise Computing Centers (DECC) should be reserved for core or sensitive data. This approach would allow the Navy to take advantage of commercial economies of scale while retaining the agility and security benefits of Cloud Computing for sensitive data.



Why This Matters to the Navy

Bechtel CTO Geir Ramleth compared his internal network costs to the costs of best-in-class providers

	Best in Class	Cost of Bechtel's Inefficiencies
Bandwidth Capacity	\$500/Mbps per month	\$10-\$15/Mbps per month
IT Efficiency	One administrator per 100 servers	One administrator per 20,000 servers
Storage Cost	\$3.75/GB per month	\$0.15/GB per month

Potential Savings are Compelling

This Bechtel Corporation case study demonstrates the positive impact of next-generation networking infrastructure on cost efficiencies in network operations. Comparisons were made using three cost performance metrics: bandwidth, IT staffing, and storage. The Panel chose Bechtel for two principal reasons: they have published internal data, which is usually difficult to obtain in open source literature; and Bechtel is a company that deploys units to remote regions where both bandwidth and connectivity are limited. This is the great challenge for Navy and Marine Corps networks – serving the user on the tactical edge. Note that Bechtel's Chief Technology Officer (CTO) is uncommonly candid in acknowledging his inefficiencies. Based on the measurement comparisons, it is clear that the potential savings are compelling.

The Bechtel network costs for bandwidth are compared to YouTube. YouTube, a social networking site, uses a network and storage provider to provide Platform as a Service (PaaS) for hosting media content to achieve considerable cost savings. Through the hosting of multiple customers within the same virtual environment and other

economies of scale, YouTube achieves substantial savings in bandwidth cost when compared to Bechtel.

Likewise, Google's network administration is significantly more efficient than Bechtel's due to virtualization and centralized administration of its very large data centers. Through the use of virtualized storage and centralized software maintenance, Amazon.com achieves 25 times the cost efficiency in storage as the legacy Bechtel hosting environment. The application of the new networking architectures allows users to reduce traditional expenditures on hardware, software, and services by paying providers on a pay-as-you-go basis. Clearly the potential savings to the Navy are compelling.

Today, there are many Navy IT inefficiencies, including:

- Tens of thousands of application versions on hundreds – possibly thousands of networks.
- Poorly supported by undermanned and sometimes untrained IT administrators.
- Thousands of server locations.
- Paying Microsoft to maintain Windows NT operating system which was released in 1993 (seven newer operating systems have been developed and launched by Microsoft since 1993).
- Poor security patch management – industry and government security experts claim that new patches are “probed” within an hour of being installed. In general, Navy IT takes months to install security patches, involving multiple program offices to validate the software code.
- From interviews with the COMTHIRDFLT N6 staff and the review of strike group post-deployment reports, the Panel learned of common IT deficiencies for the operating forces: poor training, network systems engineering, and configuration control and management; reduced bandwidth availability; “multiple systems...multiple level classification...multiple enterprises...then, finally, multiple versions...”

These are compelling reasons for the Navy to look to the means the private sector is adopting to solve many of the same problems.



Why This Matters to the Navy

“Our S&T investments must address Warfighting gaps and improve our effectiveness and efficiency.”

– 2009 CNO’s Guidance

- Information dominance is central to **National Military Strategy**
- Information dominance **requires** being near the leading edge of technology
- Being **near the leading edge** today can increase network efficiency 10s to 100s of times
- **Societal transformations** are being driven by advances in Information Technology (e.g. social networking is changing how people interact with people)
- The **new networking architectures** enable the NNE objectives as well as operational commander priorities well beyond what the current architecture can provide

Translates to greater effect at lower cost

No single slide can do justice to the effort the Panel applied to reconciling the characteristics of the next-generation networking architectures to Naval network operational imperatives. The Panel used the Naval Networking Environment goals and features as the principal guide, enhanced through a review of after-action reports from Naval strike groups and a meeting with Third Fleet N6 personnel. In particular, the new architectures accommodate the imperatives for information sharing among Naval units in joint and coalition operations. They also support the massive data processing required for timely and effective ISR – well beyond what the current architectures provide. The Panel noted that transformation across our society is being driven by advances in IT, as evidenced by the ubiquitous use of social networks by Sailors and Marines. In fact, our next generation networking environments may have Facebook-like functionality within the private cloud – although this functionality will probably not be a requirement.

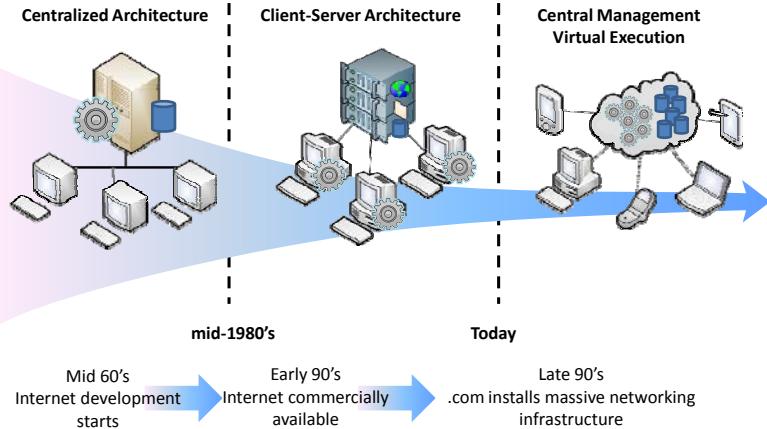
The net outcome is greater effect at lower cost, both in enterprise and tactical-edge applications. These advantages, however, will not accrue unless the Navy gets involved today.

This new architecture offers an unprecedented opportunity to combine and deliver applications, hardware, software, and data storage in revolutionary ways at significantly reduced costs. Exploiting technology can increase network efficiency 10s to 100s of times. It is expected that when the speed of information flow is fully exploited, increased efficiency, productivity, and operational effect will result – all essential enablers of our National Military Strategy.

At a more fundamental level, reconciling the Naval networking imperatives with the tremendous advances in the commercial world may be misleading. The advances occurring in the commercial world are the product of “technology push” – the efforts of a comparatively small number of visionaries – rather than “market pull.” During the technology push phase, all but a few customers resist adopting the new technology as they are unwilling to assume the risk for trying unproven approaches. However, in the commercial world, technology push can shift to market pull in a short period of time. In the market pull phase, customers adopt innovative products and processes to realize competitive advantages from the new offerings; and the new technology moves toward widespread adoption. Companies that do not make the change risk being put out of business. The lesson for the Navy is that transformational ideas may not fit into the framework of the current operational paradigms, yet still be critical to future of Naval operations.



Evolution of Networking Architectures



This slide depicts the evolution of networking technology, from the earliest days of technical computing through today's client-server architecture to the constructs emerging today that will define the next-generation architectures. This evolution is set against the development of the internet, which provides the critical networking connectivity. Each transition is caused by the inability of the previous architecture to fulfill the speed and accessibility requirements that are enabled by the next. The underpinning of the current evolution – a combination of the centralized and client-server models of the past – is virtualization. Virtualization permits the management of real-time aggregation/disaggregation of computing resources (e.g. processing power, storage) on-the-fly, as required. This opens the possibility of developing computing as a commodity, purchasing processing power and storage as we did with the centralized architecture but distributing execution to optimize the use of these resources. In the case of the client-server architecture, any increase in processing capability or storage capacity to meet peak needs requires capital expense for added capability that cannot be recovered when the pace slackens. The new architecture permits capability and capacity to be expanded or contracted on an as-needed basis.

Stage 1 – Many users share a central computer:

In the early 1960s, a single large computer ran "jobs" fed into its expensive "maw" via card readers operating on a near 24 hour schedule (minus scheduled maintenance downtime). By the mid-60s, time-sharing enabled many users to operate the computer via terminals. Time-sharing allowed for more users to get "computer time" but could do nothing to ameliorate the slow output caused by multiple users of the central computer – given the available computing "power" of the mainframes. If you wanted to share your data with the users of another computer located elsewhere, you made frequent trips to the post office with packages filled with various sorts of computer cards or tapes. By the late 60s, select users were benefiting from the ARPAnet, sharing files using an early version of File Transfer Protocol (FTP). Then, in the 70's, computers moved beyond purely technical applications and productivity computing began to take hold. Word processors, spreadsheets, and data bases (the fundamental building blocks of the productivity evolution) appeared and were steadily improved. Nevertheless, the networking architecture remained centralized, using central servers and dumb terminals.

Stage 2 – Many Personal Computer (PC) clients share a central server:

By the mid-1980s, general computing moved from an expensive central computer to a collection of inexpensive single-user PCs, operated for only six to eight hours a day.

There was still the problem of data sharing, but this was handled by replacing the old time-shared computer with a server that provided the PCs with access to data and a place to store data of interest to other users. It was natural to call those PCs the clients of the central server.

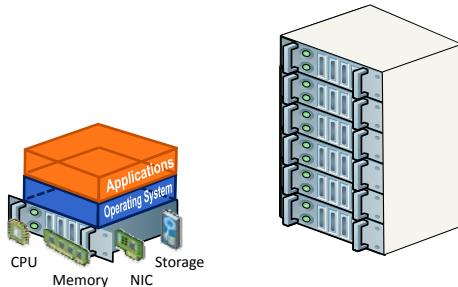
With the improvement in the speed and efficiency provided by on-board processing with central storage, the central server no longer had to be in the same room as the clients. Users could control their own computing schedule, no longer having to share a scarce resource. But, this new independence required a significant level of

professional administration to ensure routine PC maintenance, installation of programs, and security patches are properly managed. (The alternative, which is unfortunately often the situation in Naval networks, is that users work with obsolescent software on equipment that is increasingly vulnerable to malware.)

Productivity was improved by the easily accessible and varied information that was available via the Internet from the popular service providers such as America Online (AOL). By the late 1990s, use of the Internet became almost universal in most industrialized societies – with world-changing waves of information services accessible to any computer capable of operating a web browser. Today “any computer” includes the modern cell phone.

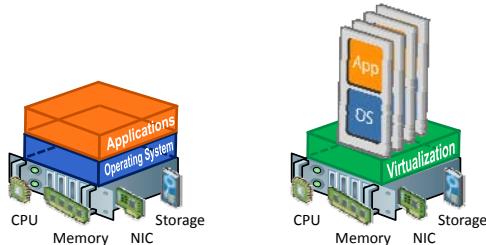
State 3 – Many portals share a collection of virtual computers running on a collection of physical processors and storage devices in a Cloud:

At first, the computing was done centrally, and users were located in proximity to the central computer. Later, processing was decentralized and performed locally by PCs, with centralized data storage. Today, with the movement toward Cloud Computing, it is neither local nor central, but distributed over multiple processors and data storage devices that can reside in multiple data centers. Users do not know *or care* which data center hardware – at whatever location – they are actually using. Users have what appears to them to be a personal, powerful, up-to-date computer; running up-to-date applications, with up-to-date security patches. To the user, the computer appears to be tightly coupled to services, databases, and Internet information suppliers.



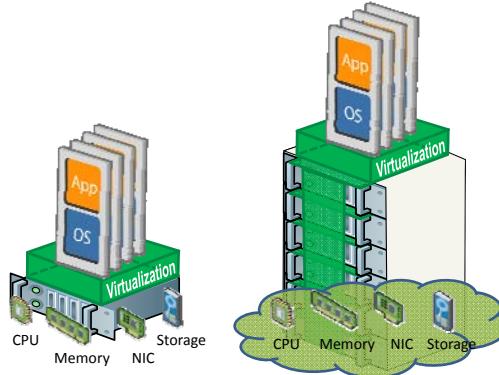
- 1-to-1 Server/OS Ratio: Operating system and applications installed on each machine.
- Configuration Management: OS and applications updated periodically, creating diversity of versions across the network.
- Security Patches: IT must support ALL of the versions!

This slide depicts how a computer is typically used today. The user “owns” a physical set of hardware and installs an operating system compatible with the hardware and the desired application software. This approach allows a high degree of tailoring to the needs of individual users. Across an organization, tailoring yields a wide diversity of hardware, operating systems, and installed applications – with each diverse system requiring its own maintenance (e.g., upgrades, security patches). The resulting complexity leads to significant administrative oversight/overhead. As an organization’s networks become larger, the IT management becomes increasingly complex.



- Flexibility: Virtualization layer allows multiple OS on an individual server.
- Standardized Configuration: Provides a uniform server environment – easier for IT to support!
- Efficiency: Allows each server to be used more efficiently, and therefore, requires fewer physical servers.

An essential element in modern IT architectures is the “virtualization layer”, which decouples the hardware and software. (Alternatively, it can be said that the virtualization layer makes the hardware compatible simultaneously with many different operating systems.) The advantage is that virtualization allows a single machine to be used as if it were many: that is, several operating systems and associated applications can be installed onto a single physical computer “box.” This has several benefits. One is that a given set of hardware can be used much more efficiently. The average utilization rate of most processors is below 10%, so loading multiple operating systems (and applications) onto a single server allows the average utilization rate to be increased by several times. Another benefit is that it allows multiple machines to be linked so they can operate as a larger, more powerful computing resource. Finally, as these centralized configurations are standardized, IT management can be optimized.



- **Scalable:** Virtualization layer supports seamless expansion of computing and storage capacity on demand.
- **Pools Resources:** Permits creation of large, shared server and storage capacity serving large and diverse user community.
- **Availability:** Load leveling virtual machines across servers provides instant recovery from failure of physical servers.

Virtualization allows the linking of hardware to provide computing resources on a massive scale – capable of supporting a large number of consumers for a variety of purposes. This is called Cloud Computing. There are three dominant versions, depending on how much of the construct is provided as a service:

Infrastructure as a Service

In this case, a vendor owns the computational hardware, exploiting economies of scale both in the purchase of the hardware and in the centralized, efficient administration by a smaller number of skilled IT administrators. Via the internet, the consumer purchases whatever resources are needed on demand. The vendor uses automated provisioning tools to bring the “virtual machines” on-line in the configuration requested and can scale-up and scale-down in cycle times of a few minutes. The user obtains benefits by avoiding capital expenses (transferring capital expenses to operational expenses), removing the need to predict resource requirements in advance, and decreasing the need to hire so many skilled IT administrators.

Platform as a Service

In this model, a vendor offers a “platform”, i.e. a standardized application-development environment. Typically, this platform is a Service Oriented Architecture (SOA) in which new applications are created from a set of previously developed, common services. The consumers, i.e., application developers, benefit by being able to concentrate on writing their applications rather than on setting up and maintaining their environment. Also, they can leverage previously developed services for their own, tailored uses. This approach enables faster innovation, as well as the iterative benefits derived from combining individual developer innovation.

Software as a Service

In this model, a vendor offers a centralized, multi-user architecture in which software applications are offered and continuously maintained for the benefit of all users (consumers). This approach ensures that all users have immediate access to the most recent versions of the software, without any IT administration burden on their part. In addition, this software may be leased for the period of use rather than purchased, allowing users to pay for what they consume instead of paying an up-front cost for hardware and software, regardless of usage.



Cloud computing is a style of computing that enables

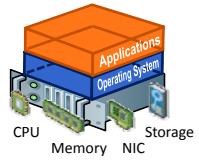
- available, convenient, on-demand network access to
- a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, services) that can be
- rapidly provisioned and released with
- minimal management effort or service provider interaction.

Adapted from the NIST Working Definition

Virtualization enables several fundamentally new ways of configuring and using hardware and software. In fact, there are almost an infinite number of variations of the three fundamental building blocks – Infrastructure as a Service, Platform as a Service, and Software as a Service. Capturing the full spectrum of variations is not possible because of the rapid and accelerating pace of innovation in this area. Thus, there is no universally agreed upon definition of “Cloud Computing.” (The various Cloud Computing definitions provided to the Panel during the fact-finding phase are listed in Appendix B.) There is, however, general agreement about a number of key elements regarding what is new and different about the Cloud. The Word Cloud in the slide above shows the dominant concepts that emerge when all the definitions the panel received from the many briefings and discussions with Cloud providers are run through “wordle.net”. (This web-based program shows a “Word Cloud” which gives greater prominence to words that appear more frequently in the source text. The larger the word, the more it recurs among the definitions.) The words that do not appear as dominant concepts are as important as the words that do, especially when considering the application of Cloud Computing to Naval missions. For example, the words security,

bandwidth, connectivity, and latency are not dominant. Yet, these network attributes are critical to Naval applications at the tactical edge. We will discuss this further with the slide that talks to Naval operational gaps.

A modified version of the National Institute of Standards and Technology (NIST) definition of Cloud Computing is shown as well. It captures the characteristics of Cloud Computing the panel considered essential.

Before:

- **7 data centers** and **35,000 sq ft** of datacenter capacity across a distributed footprint.
- **5 versions** each of **230 applications**. Upgrades and training were constant.
- **No version management.**

After:

- **3 data centers** with less than **1,000 sq ft** of datacenter space.
- **1 version** of each of the **50** most heavily used applications converted to run from a Google-like portal.
- **Centralized version management.**

Having reviewed the technological underpinnings of cloud computing, let us return to the question of how cloud computing benefits materialize. In the Bechtel Corporation example presented earlier, sample internal network figures of merit were compared with the best-in-class performers. Bechtel found they were paying far more for bandwidth and storage, and required many more IT professionals per server. The differences were factors of tens or hundreds. The conclusion drawn from their analysis was that the best-in-class performers were taking advantage of the Cloud Computing paradigm. They experienced lower storage costs through economies of scale – in effect, creating large modularized data centers. Lower bandwidth costs were realized in part by locating large data centers near large data pipelines. The reduction in IT staff resulted from moving to a centralized management of standardized servers and applications. The question for Bechtel was – given that cost savings could be achieved – should they outsource, or should they attempt to implement their own Cloud?

Bechtel made the decision to completely re-structure their computing network to take advantage of Cloud Computing rather than outsourcing. So, rather than use the

public Clouds created by companies such as Amazon and Google, they opted to create a private Cloud. They consolidated to three data centers and co-located them with telecom company operations – enabling a far more competitive environment for purchasing bandwidth. The footprint of the new data centers, filled with standardize servers providing virtualized platforms, was a fraction of the size of the previous data centers. As of this writing, the change-over is still in progress; but Bechtel is now in the process of virtualizing the top 50 software applications used by its employees – eliminating the need for most PC-hosted application software. As the Department of the Navy explores implementation strategies for Cloud Computing, early adopters like Bechtel can provide valuable insights to potential benefits and pitfalls.



Cloud Revolution = Think Scalable



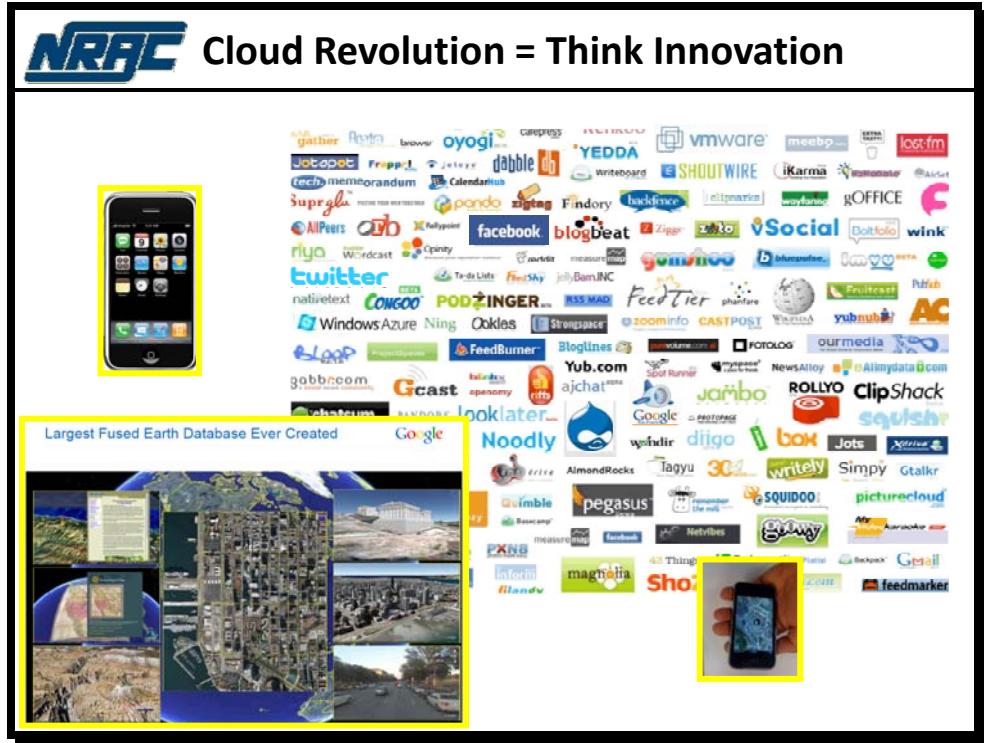
Architecture revolutions enable new paradigms in the “.com world.” For example, virtualization and automated provisioning are changing the way we think about massing hardware. In the past, when more computing power or more storage was needed, additional hardware was deployed. This quickly led to rooms full of computer racks maintained by skilled IT administrators – and these facilities required special power, cabling, air conditioning and other infrastructure.

Using the new networking architectures, the commercial world has begun to optimize data center designs. These are stand-alone enclosures that hold efficiently-packaged hardware and are located to take maximum advantage of available infrastructure as well as utility rates. The enclosures are typically manufactured directly into shipping containers, which can then be transported easily to the point of operation. Once in place, the container is connected to power and the network, and can be brought into operation almost immediately. The containers can be used in a stand-alone mode, if a limited capability is needed. For larger scale operations, modern data centers are being

constructed from hundreds of such containers. This modularity approach allows for rapid adaptation to IT needs and is an important aspect of agility in this new environment.



Cloud Revolution = Think Innovation



The Cloud revolution is real and already here. The background of the slide shows only a fraction of the companies that are providing services to the Cloud and/or consuming those services to offer new products to customers.

Google Earth is a well-known provider of Cloud-enabled capabilities. Much of this capability was developed by Google, and was possible because of their vast, “Cloud-centric” architecture that stores, retrieves, and delivers data to a huge number of users. It is also important to note that much of the content on Google Earth has been added by others (i.e., “crowd sourcing”). For instance, the Google Earth platform enables the integration of photos and documents provided by other, non-Google sources. This “crowd sourcing” is a transformation in how large groups of people share information. It is enabled by the dramatic reduction in the cost of computation, bandwidth, and storage offered by the Cloud.

Another example is the Apple iPhone. This device is celebrated for its highly innovative user interface – but there are other key aspects of interest. Almost as soon as

the iPhone was made available to customers, a number of innovative applications (known as “apps”) were available for user download. Most of these apps were not developed by Apple. They were developed – and continue to be developed – by a large number of third-party software innovators. Their add-on applications are possible because of the standard development environment for the iPhone – making the phone more valuable by continuously rolling out new capabilities. The non-linear power of innovation again is illustrated by the iPhone Google Earth download utility. This enables the mobile user to benefit from Google’s vast databases and computational resources. Obviously, the ability to harness the full power of the Cloud could be transformational for Naval Forces.



Technical Issues for Naval Implementation

Security in the Cloud

- Uniform, efficient enforcement of security standards
but
- It opens up some new and different security concerns
 - In the Virtualization Layer
 - In Software
 - In data at rest

Bandwidth & Connectivity to the Cloud

- .com implementations assume ubiquitous, high-bandwidth, continuous connectivity
but
- The Navy must deal with limited bandwidth and intermittent connectivity

While modern IT architectures offer potential benefits, the Navy has special challenges it must address – leading to some technical issues unique to Naval (or military) implementation.

Security takes on added importance in military operations, and so solutions that are sufficient for “.com” may not pass muster for Naval applications. For instance, virtualization is a key enabler of modern IT architecture – and it is also potentially a key vulnerability, because the virtualization layer controls all access to hardware and software. (In the commercial world this function is called the “hypervisor.”) The military must ensure that this virtualization layer is safe from cyber attack.

Also, the strength of the Cloud is that it holds data and applications separate from the hardware. The advantages of this approach were discussed earlier, but it is important to recognize that this loss of control in where data is stored and where applications are

actually being run brings additional security challenges that must be addressed. As more users move to the Cloud – the banking community for example – they will require this gap to be closed. But the military requirements will most probably never be satisfied by a “.com” security solution.

Assuring user connectivity is a fundamental, underlying assumption of Cloud Computing. This assumption is a good one for most terrestrial applications, but it does not necessarily convey connectivity requirements for Naval operations. Operations ashore typically have essentially unlimited connectivity, but operations afloat have much more limited bandwidth – with dismounted operations (i.e., Marine infantry) having even less. The Navy will not be able to rely on commercial vendors to optimize Cloud solutions for highly limited and intermittent bandwidth. As a single example of the looming challenge, the web-based data exchange standard is Extensible Markup Language (XML), a format that requires extensive bandwidth. The Navy’s narrow pipes at the tactical edge will not easily support this format. At times connectivity may be lost completely, and this brings a new set of challenges in how to re-integrate previously disconnected parts of the Cloud, including how to re-establish network trust.

The very innovation and speed-to-market attributes of Service Oriented Architecture (SOA), enabled by the Cloud, present inherent challenges for the Navy. It must adopt a new testing model to ensure that it keeps pace with the new paradigm in which developers can rapidly reconfigure existing services to provide new capability. DISA is undertaking a similar approach to driving innovation in converting from its legacy Global Command and Control System (GCCS) to its modern, SOA-based Network Enabled Command and Control (NECC).

In addition, decomposing an application into a series of services may result in unexpected time delays. Latency is unimportant in most commercial applications but is much more critical to military operations. Recently, the commercial sector has become more concerned with latency (e.g., for streaming video) despite generally large bandwidths available. It appears they are seriously developing new protocols to mitigate

this problem. This development could help the Naval implementation but will most probably not resolve the issue completely.

Finally, standards are of concern – especially in the area of SOA. If the Navy develops capabilities using an application development environment provided by a commercial vendor, the Navy could become hostage to that vendor. Similarly, a vendor could make it difficult for the Navy to extract data from commercially hosted data centers. Thus, the Navy will need to carefully weigh concerns about standards against the simplicity of commercially provided Cloud capacity.



Finding I: Transformational

Finding:

1. Cloud computing technology has the potential for transformational benefits to Naval networks (not a fad).

Recommendations:

- 1-A. Acquisitions: Ensure future DoN acquisitions consider and, as appropriate, leverage the benefits of cloud computing.
- 1-B. Pay-as-you-go: Develop long-term procurement strategies for purchasing on demand computing capacity.
- 1-C. Metrics: Standard cloud computing models and performance metrics should be developed to assist in the design, monitoring and contracting of systems.
- 1-D. Pilot: Establish cloud computing pilot program(s) to explore the key metrics, benefits and issues.
- 1-E. Standards: Must enter the standards conversation with the commercial community to represent unique Naval needs.

Commercial industry is developing and adopting Cloud Computing at an astonishing rate. A sea change in networking has arrived, and major improvements in network efficiency are possible. The new networking architectures will open the opportunity to improve efficiency and effectiveness in the ways we develop and employ hardware, software, and IT personnel. Configuration management in our networks can now reach levels that were heretofore unattainable. And significant cost savings and cost avoidance will be available through different financial vehicles that are more compatible with Naval budgetary needs.

Accordingly, the Panel makes the five recommendations shown in the slide. The first (1-A) is a long-term recommendation which seems almost trite. However, it reflects the reality that the most difficult step is always the first. If the Navy is to take advantage of this opportunity, it must take the new networking architectures on board as a future imperative and begin looking to Cloud capabilities to improve its efficiency and effectiveness today. The next three (1-B, 1-C, and 1-D) are very closely related and support the first recommendation directly. They speak to the need for network simulation

models that are developed and validated with empirical data obtained from field experiments. The Panel suggests that the place to begin experimenting is not at the tactical edge but rather in land-based administrative functions that require non-operational information sharing with non-military personnel, e.g., dependents. Family services functions might be a good place to start. Since some of the data exchanged may be personal, running pilot programs will afford the opportunity to explore how the Navy and Marine Corps can handle private information in the public cloud. The lessons learned will support the metrics needed to plan adoption and implementation as well as to monitor performance. The last recommendation (1-E) talks to the need for the Navy to be involved in the standards development discussions with commercial developers, commercial users, and government entities (e.g., NIST). The purpose would be to remain abreast of developments as well as to influence the direction of development when necessary. This will permit the Navy and Marine Corps to take maximum advantage of the COTS products without a great deal of modification.



Finding II: Security

Finding:

2. Trust in the security of cloud technologies; i.e., confidentiality and integrity, is the greatest challenge to cloud utilization.

Recommendations:

- 2-A. Research Areas: Track research and fund activities to fill Navy specific gaps:
 - Trusted (formally verified) virtualization layer
 - Data-at-Rest in the cloud
 - Secure cloud applications
- 2-B. Confidentiality/Integrity: Develop strategies, technologies, and protocols to enable Naval forces to fight through loss-of-trust events and to rapidly restore trust and integrity of cloud operation. Future Naval war games should test these strategies, technologies, and protocols.

Data and communications security, while important in commerce, assumes paramount importance in a warfighting context. The Cloud paradigm presents new challenges to establishing and certifying cyber security. These new challenges are unlikely to be met by commercial vendors; hence, the military research enterprise must address them to make Cloud Computing useful to the warfighter. It will be a challenge for the military to engage the appropriate technical communities to ensure that its needs are met. Unlike cryptology, where there is a historical tie between the national security establishment and the academic community, the leading IT practitioners that are driving the Cloud Computing movement – in both industry and academia – may be familiar with military requirements but not necessarily incentivized to address them.

The Panel recommends that the following critical Cloud Computing research challenges be addressed:

Security of Virtualization Layer: The core technology enabling Cloud Computing is the virtualization layer (VL) of the architecture. Standing between the computing

hardware and the operating systems and applications, it is the main reason that Cloud Computing represents a paradigm shift from client-server architectures. It represents a potential single point of failure with regard to information security. The VL, by definition, could have access to every virtual machine operating in a Cloud, with the associated processors, storage, Operating System (OS), and applications operating on them. This universal access underpins the ability to push patches in an extremely efficient way. On the other hand, if a virtualization layer contains a vulnerability, exploitation of that vulnerability by an adversary would allow compromise of every process operating in the Cloud, as well as all data in the Cloud.

Vendors recognize this dual-edged problem. Several vendors explicitly promote their VLs as compact and tightly engineered, suggesting that their VL presents a harder target for potential compromise relative to their competitors' offerings. However, most vendors do not publish source code for their VL, nor do they submit them for code analysis. The primary differentiation between VLs is size and the process discipline exercised by the vendors' software developers. It goes without saying, that off-shore code development represents a particular threat in the context of VLs. Finally, it should be noted that software "bloat" (i.e., when software has a larger than necessary footprint with many unnecessary features that are not used by end users) will tend to increase the vulnerability of VLs, as well as obscure the vulnerabilities. The Panel estimates that the most tightly engineered VLs commercially available today approximately conform to only EAL-4 (on a one to seven scale) of the NSA/NIST National Information Assurance Program (NIAP). EAL-4 is "Evaluation Assurance Level 4: Methodically Designed, Tested and Reviewed". EAL-7 is "Formally Verified Design and Tested".

Two potential remedies to this aspect of Cloud security exist. First, a VL developed using a standards-based open-source approach would likely have fewer and more rapidly disclosed vulnerabilities than a proprietary offering. Strong U.S. Government engagement in the VL open-source community probably provides the strongest near-term support to VL security.

In the longer term, there is another way to ensure security of the Cloud: formal verification of the virtualization layer. Formal verification of software has been studied at least since the 1960s. Based on mathematical theorem proof, formal verification can assure that entire classes of vulnerabilities are absent from a piece of software (e.g., buffer overflow). The primary difficulty associated with formal verification is that it has remained computationally intractable for practical-sized pieces of software. However, the so-far-inexorable advance in computational power driven by Moore’s law has enabled the formal verification of larger and larger programs. Naval research should attempt to establish whether a basic VL code might be amenable to formal verification in the foreseeable future. This possibility also motivates keeping the VL as small as possible. Intermediate results between EAL Level 4 and EAL Level 7 would be useful enhancements to Cloud security.

Data at Rest: A second area where information security challenges require a new perspective is the vulnerability of data at rest. First, the simple encryption of data at rest weakens many of the benefits of collecting data in a central repository. For example, simple search over encrypted data is not currently practical, but remains an active research topic. The indexing of encrypted data in a Cloud, at best, becomes the responsibility of the user rather than the Cloud provider. More elaborate uses of data, such as mash-ups (e.g., using cartographic data from Google Maps to add location information to real estate data – not originally provided by either source) are not even well-defined in the context of encrypted data.

Cloud vendors assert that their proprietary processes provide significant security through segmenting customer data, and redundantly storing it in small fragments over geographically distributed Cloud storage servers (i.e., the process known as “sharding”). However, the proprietary algorithms that vendors use to shard data are attractive targets for reverse engineering. In conjunction with VL vulnerabilities, data at rest would be vulnerable.

Operating System and Application Security: Even in the context of a secure VL, it is still entirely possible (in fact, typical) for the supported OS and applications to exhibit vulnerabilities. Microsoft Windows OS operating on top of a secure VL will remain just as vulnerable as on a stand-alone server. The VL may provide some useful form of containment, by restricting communications between different virtual machines (VMs) to certain approved modalities. Further, intentional collapse of VMs between jobs (with re-instantiation from known-good sources) will also tend to limit propagation of vulnerabilities, exploits, and malware. However, vulnerable application and OS software will remain a central problem for information assurance in commercial and military systems for the foreseeable future.

The Panel received one briefing from a leading computer scientist indicating that a fundamental reengineering of IT infrastructure based on functional programming offers substantial hope for secure applications and infrastructure. (His research is being supported by the Air Force Research Laboratory.) This reengineering represents a huge undertaking, and even stipulating that the conjecture is true, could not be realized in the near-term. However, this research area is indicative of the trends in information technology of which the DOD *must* retain awareness.

In the area of user confidentiality and integrity, the Panel recommends that the Navy develop strategies, technologies, and protocols to enable Naval forces to fight through loss-of-trust events and to rapidly restore trust and integrity of cloud operation. Naval war-games should test these strategies, technologies, and protocols.



Finding III: Intermittent Operations

Finding:

3. Naval forces afloat will remain disadvantaged by intermittent links with large latency and marginal bandwidth. This may affect the capabilities of cloud computing-based systems that depend upon more reliable links.

Recommendations:

- 3-A. Continuity: Develop technologies to ensure continuity of cloud operations in the face of failed communication links (e.g. between shore and afloat components).
- 3-B. Synchronization: Research ways for cloud synchronization over intermittent/low-bandwidth/mobile channels
- 3-C. Redundancy: Research and develop high bandwidth and multiple redundancy links (e.g., the DARPA ORCA program).

U.S. warfighting communications requirements are designed to support information exchange enabling C4ISR (command and control, communications, and computers; intelligence, surveillance, and reconnaissance), battle management, and combat support information. Today's Naval combatants are equipped with limited communications bandwidth to support voice and data requirements. This legacy is due to the difficulty of communicating while operating over long distances, in challenging weather and in potential jamming environments. Typical communication bandwidths provided on an Aegis cruiser are: Link 11 – 2.2 kbps, JTIDS/Link 16 – 56 kbps to 112 kbps (depending on the degree of jam resistance needed by current operations). Even the special-purpose Cooperative Engagement Capability (CEC) link is only 50 kbps all-to-all (for typical network sizes) or 2 Mbps for one-to-one. The SIPRNet, when available via UHFSATCOM, is only 256 kbps.

However, there are emerging DOD systems that will expand the current capabilities for both strategic and tactical users. DOD is populating a constellation of

Wideband Global SATCOM satellites while upgrading legacy communications satellite systems: the Navy UHF Follow-on (UHF F/O) is evolving to the Mobile User Objective System (MUOS); the Air Force Milstar system is evolving to the Advanced Extremely High Frequency (AEHF) satellite communications system.

The Wideband Global SATCOM (WGS) architecture will provide enhanced communications capabilities for U.S. and allied forces deployed around the world. WGS also augments the current Ka-band Global Broadcast Service (on UHF F/O satellites) by providing additional information broadcast capabilities. The combination of the Wideband Global Satellites, current Defense Satellite Communications System (DSCS) satellites, Global Broadcast Service (GBS) payloads, wideband payload and platform control assets, and associated earth terminals comprise the Interim Wideband System (IWS). The IWS supports 24/7 wideband satellite services to tactical and selected fixed infrastructure users. Limited protected services will be provided under conditions of stress to users employing anti-jam terrestrial modems.

Accordingly, the Panel recommends the following for intermittent operations:

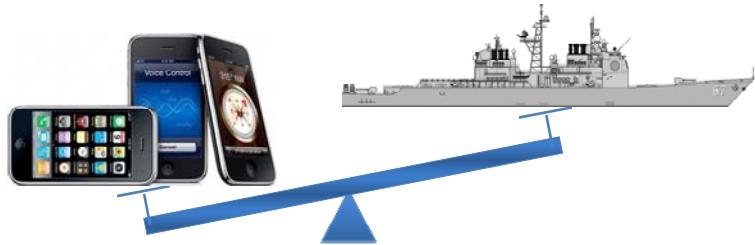
Ensuring Continuity – Naval Private Clouds afloat can operate in concert with the larger Naval Private Cloud ashore. These shore-based private Clouds can also access commercial Clouds as needed in a dynamic way. Naval combatants must maintain readiness and operational capability even in the event of a loss of the connectivity between the ship-based private Cloud and the shore-based private Cloud. Research on how to best mirror the information databases that reside in the ashore private and commercial Clouds should be undertaken. Ashore Clouds will need to be exploited for computational tasks thereby load-leveling the processing needs of the ships at sea.

Ensuring Synchronization – Line of sight, Radio Frequency (RF) data link losses aboard ship may be caused by loss of track, weather phenomena, or own-ship blockage. The ship environment for communications has always been challenging – line of sight links must be compensated for shipboard motion, including blockage by the ship's super-

structure. Ships that report into or depart from an operating task force will also require communications synchronization. When a link is broken and then reestablished, the ship-based Clouds will begin to refresh and update their local databases. Care must be taken that this refresh process occurs in a prioritized, well conceived fashion so as not interfere with other real-time data link traffic. Research and development of tools and technologies to manage these re-synchronization tasks is needed – and ONR should take the lead in developing such programs.

Ensuring Redundancy – The Department should strive to improve the wireless bandwidth available to ships at sea. Near-term enhancements can be achieved by implementing satellite receivers for WGS and AEHF. Beyond increasing individual link bandwidth, research in creating redundant wireless links to both satellites and airborne relays should be considered. The Broad Area Maritime Surveillance (BAMS) UAV might serve as a future high altitude relay platform. ONR should also explore redundancy within a channel by leveraging both optical and RF carrier frequencies. For example, ORCA (Optical RF Communication Adjunct) is a DARPA optical and RF communications project to provide a high data rate gateway network capability to warfighters. It will include airborne nodes; on-the-move and stationary ground vehicles; and Global Information Grid (GIG) Points of Presence. ONR could extend this technology development to include links to Naval ships and aircraft at sea.

Bandwidth Challenge to DoN Use of Cloud Architecture



**Connectivity pervasive for .com applications
Achieving this for disadvantaged user is nontrivial**

As has been stated, the Naval Force afloat is disadvantaged with intermittent and marginal bandwidth links. This may affect the capabilities of Cloud Computing that assumes robust, high-bandwidth links. The Panel has made several recommendations concerning continuity, synchronization and redundancy aimed at ameliorating this condition.

It is worth noting the magnitude of this problem, given the pervasive connectivity that we have come to expect in our homes, offices and vehicles. In short, almost all of us, with the possible exception of those in certain scientific disciplines, have more bandwidth than we need – around-the-clock, wherever we happen to be. A mobile phone user with the newest iPhone (model 3G S) operating on a HSPA (High Speed Packet Access) network can support up to 7.2 Mbps. Current AT&T networks in the U.S. can support up to an effective throughput of 1.7 Mbps for the iPhone 3G S; with maximum throughput of 700 kbps (0.684 Mbps) for an average iPhone.

The typical bandwidth allocated for a modern Aegis Class cruiser is approximately 2.048 Mbps – an E1-size data link. For this bandwidth comparison, the Panel used the Aegis cruiser USS Lake Champlain (CG-57) with its entire major communications links available, with some footnotes. Link 16 was included with full anti-jam capability enabled. Also included were the SIPRNet, with a typical bandwidth allocation for classified communications, and ship-to-ship Cooperative Engagement Capability with the average bandwidth of a typical network in a battle group.

What is quite surprising is that the total bandwidth of all these major shipboard circuits is almost directly comparable to that of a single iPhone. The slide above suggests that anyone with three iPhones in hand would own significantly more bandwidth than the Commanding Officer of a modern cruiser!

The Navy does not deliberately constrain its ships or deliberately disadvantage them in terms of their ability to conduct missions. But achieving increased bandwidth for ships at sea is not a trivial engineering problem. And there are other communities within the naval service whose connectivity is even more limited than the surface fleet. The submarine force suffers from intermittent connectivity, while expeditionary forces have only limited connectivity in remote areas. The bandwidth challenge is very real, and its solution is essential for deployed forces to benefit from the fast evolving benefits of new networking architectures.



Actions

- **ASN RDA**

- 1-A. **Acquisitions:** Ensure future DON acquisitions consider and, as appropriate, leverage the benefits of cloud computing.

- **DON CIO**

- 1-E. **Standards:** Must enter the standards conversation with the commercial community to represent unique Naval needs.

- **N6**

- 1-D. **Pilot:** Establish cloud computing pilot program(s) to explore the key metrics, benefits and issues.

- **NNWC**

- 1-C. **Metrics:** Standard cloud computing models and performance metrics should be developed to assist in the design, monitoring and contracting of systems.

- 2-B. **Strategy:** Develop strategies to enable Naval forces to fight through loss-of-trust events and to rapidly restore trust and integrity of cloud operation. Future Naval war games should test these strategies.

- **PEO EIS**

- 1-B. **Pay-as-you-go:** Develop long-term procurement strategies for purchasing on-demand computing capacity .



Actions

- **PEO C4I**

- 2-B. **Technologies & Protocols:** Develop technologies and protocols to enable Naval forces to fight through loss-of-trust events and to rapidly restore trust and integrity of cloud operation.

- 3-A. **Continuity:** Develop technologies to ensure continuity of cloud operations in the face of failed communication links (e.g., between shore and afloat components).

- **CNR**

- 2-A. **Research areas:** Track research and fund activities to fill Navy specific gaps:

- Trust (formally verified) virtualization layer
- Data-at-rest in the cloud
- Secure cloud applications

- 3-B. **Synchronization:** Research ways for cloud synchronization over intermittent/low-bandwidth/mobile channels.

- 3-C. **Redundancy:** Research and develop high bandwidth and multiple redundancy links (e.g. the DARPA ORCA program).

ASN RDA, Action 1-A:

Acquisitions: Ensure future DON acquisitions consider and, as appropriate, leverage the benefits of cloud computing.

The state-of-the-art in cloud computing is capable of supporting some computing requirements for the Navy today, and acquisition programs should be structured to be amenable to cloud computing solutions (see 1-D and 1-B below).

DON CIO, Action 1-E:

Standards: Must enter the standards conversation with the commercial community to represent unique Naval needs.

Cloud computing today is comprised of a variety of architectures supported by often incompatible component technologies and is an object of fierce competition. At present there are no cloud computing standards, although industry has initiated the discussion of standards. The Navy should be represented as the various stakeholders maneuver to ensure that emerging standards support Naval needs. Although the visibility of such engagement to the rest of the Navy might be low – the impact could be very high.

N6, Action 1-D:

Pilot: Establish cloud computing pilot program(s) to explore the key metrics, benefits and issues.

How should the performance of a cloud computing solution be measured? What are appropriate requirements to specify in cloud computing procurements? How do various requirements trade off against each other and what do they cost? To satisfy this need, the Panel recommends N6 establish pilot programs that implement cloud computing for discrete Navy needs, and that NNWC develop models and metrics targeted at supporting procurement processes.

NNWC, Actions 1-C, 2-B:

Metrics: Standard cloud computing models and performance metrics should be developed to assist in the design, monitoring and contracting of systems.

The ability to model cloud computing systems and the development of design “rules of thumb” will be key to an efficient and reliable procurement process. This activity will be informed by the pilot programs implemented by N6.

Strategy: Develop strategies to enable Naval forces to fight through loss of-trust events and to rapidly restore trust and integrity of cloud operation. Future Naval war games should test these strategies.

Cyber-attacks on Navy systems can have an impact that extends far beyond the scope of the original incident. A key problem is re-establishing trust in the system, including ensuring the consequences of the initial attack are fully understood and mitigated, and that the system is secure against future attacks of the same kind. This is challenging even for existing network architectures. Since the security aspects of various cloud computing architectures are poorly understood, the use of war games to explore this challenge will provide important insights. This activity needs to be closely coupled with the PEO C4I action 2-B which focuses on the development of the technologies and protocols to enable Naval forces to fight through loss-of-trust events.

PEO EIS, Action 1-B:

Develop long-term procurement strategies for purchasing On-demand computing capacity.

In the cloud computing model, computational needs may be satisfied as a purchase of services (OPEX) rather than as a purchase of hardware and commitment to

internal IT staffing (CAPEX). For some Navy needs, the commercial cloud will provide satisfactory solutions; consequently, procurement strategies need to accommodate the new ways in which computing power and mass storage is purchased.

PEO C4I, Actions 2-B, 3-A:

Technologies & Protocols: Develop technologies and protocols to enable Naval forces to fight through loss-of-trust events and to rapidly restore trust and integrity of cloud operation.

This activity complements that described under NNWC (2-B).

Continuity: Develop technologies to ensure continuity of cloud operations in the face of failed or failing communication links (e.g., between shore and afloat components).

With few exceptions, existing cloud computing implementations function within a high-bandwidth, continuously available terrestrial network. For cloud computing to benefit afloat and remote Naval forces, it must be capable of operating even when communication links are slow and unreliable.

CNR, Actions 2-A, 3-B, 3-C:

Research areas: Track research and fund activities to fill Navy specific gaps:

- *Trust (formally verified) virtualization layer,*
- *Data-at-rest in the cloud,*
- *Secure cloud applications.*

The Navy and Marine Corps have security needs that require technology unlikely to be developed in the commercial domain. ONR should identify these needs and develop research tracks to address these gaps. The nature of this research must be strong, shaped by the recognition that investment by the National Science Foundation (NSF) in networking R&D is several times greater than the DOD investment. Commercial IT R&D

spending significantly outstrips all government investment. This trend in investment drives human capital and makes it essential that DOD and the Navy establish ways to keep abreast of emerging technology in networking and computation – recognizing that its own personnel will not likely be dominant players in that development. What the Navy does have, in abundance, are difficult problems that can drive research. Appropriate engagement with academic and industry researchers may allow the government to anticipate future developments with transformational potential, such as Cloud Computing. Examples include formal verification of the virtualization layer, security of data-at-rest in the cloud, and secure cloud applications.

Synchronization: Research ways for cloud synchronization over intermittent/low-bandwidth/mobile channels.

Naval forces must often function despite low speed, unreliable communication links. Cloud implementation strategies that provide the required level-of-service despite such communication challenges should be a high priority for research.

Redundancy: Research and develop high bandwidth and multiple redundancy links (e.g., the DARPA ORCA program).

The importance of communications to cloud computing encourages the development of high-bandwidth redundant communication links. As discussed previously, ORCA is a laser-augmented system that is designed to enhance bandwidth between air and ground units. Monitoring the development of this technology and, if possible, adapting it to the maritime environment is an important first step in a bandwidth enhancement research program.



Take-Aways

- Cloud Computing: the next big step in networking architecture
- Engage the cloud community to ensure Navy needs are incorporated into evolving standards.
- Establish cloud pilot project(s) for non-combat services.
- Focus research and development efforts on:
 - Securing the virtualization layer
 - Develop data links that enable cloud architectures
 - Cloud performance models to analyze network performance in various conditions

This page intentionally left blank

Appendix A

Terms of Reference Document

Future Naval Use of COTS Networking Infrastructure NRAC Summer Study 2009

Objective

Study the Navy's use of commercial architectures, software, and hardware in its networks. This study will examine the related emerging networking approaches under development in the commercial world as well as the development and operational practices associated with them, and suggest strategies for leveraging ongoing civilian investment for Navy needs, including identifying S&T to adapt the commercial technologies to Navy operational and administrative requirements where necessary.

Background

Naval operations of all types, ranging from routine administrative activities to delivery of weapons on time and on target, increasingly depend on networked computer systems. A directive issued in early December 2008 requires the Navy to integrate existing Navy IT networks, services, and systems into a system currently labeled the Next Generation Enterprise Network (NGen). The NGen project will also support the broader Global Information Grid and Net-Centric Enterprise Services; the CIO of the Navy has even used the term "weapons system" to describe the NGen. It is reasonable to assume (certainly from the standpoint of cost) that the network architecture upon which NGen is built will make the maximum use of commercial developments. The adoption of technology that originated outside the control of the DOD for networking Navy/Marine Corps systems and people can create substantial challenges and risks.

Specific Tasking

- Compare and contrast the Navy/Marine Corps needs to maintain secure network functionality in the face of dramatically changing bandwidth availability, uncertain connectivity, and large latencies with capabilities offered by existing and emerging commercial technology.
- Explore how the integrity of Navy/Marine Corps networks can be assured with commercially developed components, e.g. when personnel developing commercial software will most certainly include non-US citizens.
- Review current Navy S&T and develop a set of actionable recommendations for new investments as well as changes to current investments that must be made by the Navy in its S&T portfolio to exploit commercial networking technologies.

This page intentionally left blank

Appendix B

Sample Definitions of Cloud Computing

Organization	Definition
Symantec	“a style of computing where scalable and elastic IT-enabled capabilities are provided “as a service” to external customers using Internet technologies” ... “Where the consumers of the services need only care about what the service does for them, not how it is implemented”
Google	“massive data centers, purpose built hardware, software platform of internet scale; Cloud Computing benefits: radically lowers the cost, much faster application development, happier end users; agile development changes timeline of product: short term is in weeks, long term is in months”
Microsoft	“Cloud Computing means using a remote data center to manage, scalable, reliable, on-demand access to applications”
Amazon.com	“new model for resource delivery; IaaS - Infrastructure As A Service; programmable data center; over the internet; flexible; on-demand; pay-as-you-go”
IBM	“standardization of deployment, integration and access to data, services, applications and capabilities (e.g. virtual machines, calling interfaces). This creates the platform automated, remote management and deployment of services, infrastructure updates; automated policy enforcement; scalable independent of operations staff; automated, prioritized capacity management to meet Service Level Agreement (SLA) requirements as workload, performance needs dictate; role-based access to services and data, required to share the infrastructure with varied constituents”
U.S. General Services Administration	“Cloud Computing is a pay-per-use model for enabling available, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, services) that can be rapidly provisioned and released

	<p>with minimal management effort or service provider interaction. This Cloud model promotes availability and is comprised of five key characteristics, three delivery models, and four deployment models. Key characteristics: On-demand self-service, ubiquitous network access, location independent resource pooling, rapid elasticity, and pay per use. Delivery models: Cloud Software as a Service (SaaS), Cloud Platform as a Service (PaaS), and Cloud Infrastructure as a Service (IaaS). Deployment models: Private Cloud, Community Cloud, Public Cloud, and Hybrid Cloud”</p>
Open Data Group	“Clouds provide on-demand resources or services over a network, often the Internet, with the scale and reliability of a data center”
University of California, Berkeley	“Cloud Computing refers to both the applications delivered as services over the Internet and the hardware and systems software in the datacenters that provide those services. The services themselves have long been referred to as Software as a Service (SaaS). The datacenter hardware and software is what we will call a Cloud. When a Cloud is made available in a pay-as-you-go manner to the general public, we call it a Public Cloud; the service being sold is Utility Computing. We use the term Private Cloud to refer to internal datacenters of a business or other organization, not made available to the general public. Thus, Cloud Computing is the sum of SaaS and Utility Computing, but does not include Private Clouds. People can be users or providers of SaaS, or users or providers of Utility Computing. We focus on SaaS Providers (Cloud Users) and Cloud Providers, which have received less attention than SaaS Users. From a hardware point of view, three aspects are new in Cloud Computing. The illusion of infinite computing resources available on demand, thereby eliminating the need for Cloud Computing users to plan far ahead for provisioning. The elimination of an up-front commitment by Cloud users, thereby allowing companies to start small and increase hardware resources only when there is an increase in their needs. The ability to pay for use of computing resources on a short-term basis as needed (e.g., processors by the hour and storage by the day) and release them as needed, thereby rewarding conservation by

	letting machines and storage go when they are no longer useful”
McKinsey & Company	“Clouds are hardware-based services offering compute, network and storage capacity where hardware management is highly abstracted from the buyer; buyers incur infrastructure costs as variable operations expense; and infrastructure capacity is highly elastic”
VMware	“lightweight entry/exit service acquisition model, consumption based pricing, accessible using standard internet protocols, elastic, improved economics due to shared infrastructure”

This page intentionally left blank

Appendix C

NRAC Panel Membership	
RADM John T. Tozzi, USCG (Ret.) (Chair) L-3 Communications	Dr. John C. Sommerer Johns Hopkins Applied Physics Laboratory
Dr. James Bellingham (Co-Chair) Monterey Bay Aquarium Research Institute	Professor Patrick H. Winston Massachusetts Institute of Technology, CSAIL
Dr. Amy E. Alving Science Applications International Corporation	Dr. David A. Whelan The Boeing Company
VADM Bill Bowes, USN (Ret) Consultant	Mr. James L. Wolbarsht DEFCON®, Inc.
RADM Daniel R. Bowler, USN (Ret.) Lockheed Martin Corporation	Mr. Ryan Gunst ONR Science Advisor to N6 (Chief of Panel Executive Secretariat)
RADM Erroll Brown, USCG (Ret.) International Business Machines Corporation	LT Josh O'Sullivan, USN Executive Secretary
Dr. Mark G. Mykityshyn White Oak Group	

This page intentionally left blank

Appendix D

Fact-Finding Contributors

Contributor	Organization
Mr. Jeffrey Barr	Web Services Evangelist, Amazon
Mr. Justin Burks	Web Services Alliance Manager, Amazon
Mr. Mike Culver	Web Services Evangelist, Amazon
Mr. Paul Horvath	Solutions Architect, Amazon
Mr. Stephen Schmidt	General Manager, Enterprise & Federal AWS, Amazon
Mr. Matt Tavis	Solutions Architect, Amazon
CAPT Nancy King-Williams	US Third Fleet, N6
CAPT Jeff Saunders	US Third Fleet, N6
Mr. Tim Dowd	CISCO
Mr. Stephen Orr	CISCO
Mr. Bruce McConnell	CSIS
Mr. Eric Gunderson	President, Development Seed
Mr. Richard Hale	Chief, Information Assurance Executive, DISA
Mr. Dave Mihelcic	CTO, DISA
Mr. Dave Baciocco	CTO, Ericsson
Mr. Kevin LaMontagne	Gartner
Mr. Robert Mason	Gartner
Mr. Jason Cain	Google Earth Enterprise Sales Engineer
Mr. Dylan Lorimer	Strategic Partner Manager, Geo Content Partnerships, Google

Mr. Rajen Sheth	Senior Product Manager, Google Applications
Mr. Graham Spencer	Applications Engineer, Google
Mr. Mark Wheeler	Google Earth/Maps Enterprise, Google
Mr. Jim Young	DOD Sales Manager, Google
Mr. Bryan Atwood	Product Manager, Google Earth Enterprise
Ms. Casey Coleman	CIO, GSA
Mr. Jay Magnino	Navy Client Manager, IBM
Mr. Lawrence Hale	CTO & PM - IT Infrastructure Line of Business, GSA
Dr. Christopher Codella	Associate Director of Technical Strategy, IBM
Mr. Alex Morrow	IBM Fellow, IBM
Mr. Herb Kelsey	Deputy CTO - Cyber Security, IBM Federal
Mr. Jeff Havens	Architect, Windows Azure – Enterprise Strategy, Microsoft
Mr. Yousef Khalidi	Distinguished Engineer, Cloud Infrastructure Services, Microsoft
Mr. Brian LaMacchia	Software Architect, Microsoft
Mr. Jeff Mendenhall	Dir Business Development-Data Center Futures, Microsoft
Mr. Dan Reed	Managing Director, Scalable and Multicore Systems, Microsoft
Mr. Dan Fay	Dir External Research for Earth, Energy, & Environment, Microsoft
Dr. Dennis Gannon	Dir, Applications for Cloud Computing Futures, Microsoft
Dr. Eric Horvitz (former NRAC member)	Principal Researcher and Research Area Manager, Microsoft
Ms. Kristin Lauter	Principal Researcher, Cryptography Research Group, Microsoft

Mr. David Aucsmith	Sr. Director, Inst. for Advanced Technology in Government, Microsoft
Mr. Brad Mercer	Chief Architect Naval C4I Systems, MITRE
Mr. Geoff Raines	Principal Software Systems Engineer, MITRE
Dr. John Gauss	NGEN SPO
Mr. Timothy Grance	Program Mgr Cyber & Network Security Program, NIST
Mr. Richard Mathews	Director, Information Assurance Research Laboratory, NSA
Mr. Ryan Gunst	Science Advisor, OPNAV N6
Dr. Bobby Junker	Head, C4ISR, ONR
Dr. Das Santu	Program Officer, Communications and Networking, ONR
RDML David Simpson	OPNAV N6N
Mr. Robert Grossman	Founder and Managing Partner, Open Data Group
Mr. John McDonnell	Asst. PEO -Science & Technology, PEO C4I
Mr. Charlie Suggs	Technical Director, PEO C4I
Mr. Gary Shaffer	Deputy Technical Director/Chief Engineer for SOA, PMW 160
Mr. Allen Armstrong	APM, PMW-140/DJC2 JPO
Mr. Rob Wolborsky	Program Manager PMW-160
Dr. Anupam N. Shah	Chief Scientist/Engineer, Enterprise & Mission Solutions, SAIC
Dr. Frank Perry	CTO, Defense Solutions Group, SAIC
Mr. Bill Vass	President and COO, Sun Microsystems Federal
Mr. Mark Bregman	CTO, Symantec

Mr. Joe Pasqua	VP Research – Symantec Research Labs, Symantec
Dr. Zulfikar Ramzan	Tech Director, Security Tech and Response, Symantec
Mr. Al Kohnle	USFFC, MOC Project Team
CAPT Mark Lane	USFFC, MOC Project Team
Mr. Steve Ebbets	Senior Systems Engineer, USN/USMC/DOD, VMware
Ms. Melissa Palmer	Strategic Account Manager, USN/USMC, VMware
Dr. Marv Langston	Former CIO, Department of the Navy
Dr. Bruce Wald	Former Director, Space and Communications, NRL

Appendix E

Acronyms

ADNS	Automated Digital Network System
ASN(RDA)	Assistant Secretary of the Navy for Research, Development, Acquisition
BAMS	Broad Area Maritime Surveillance
BLII	Base Level Information Infrastructure
CANES	Consolidated Afloat Network Enterprise Services
CAPEX	Capital Expenditure
CARS	Cyber Asset and Reduction
CEC	Cooperative Engagement Capability
CND	Computer Network Defense
COTS	Commercial Off-the-Shelf
CTO	Chief Technology Officer
DON	Department of the Navy
EHF	Extremely High Frequency
FTP	File Transfer Protocol
GCCS	Global Command and Control System
GCCS-M	Global Command and Control System - Maritime
GIG	Global Information Grid
HSPA	High Speed Packet Access
IaaS	Infrastructure as a Service
IDA	Institute for Defense Analyses
IP	Internet Protocol
ISNS	Integrated Shipboard Network System
IT	Information Technology
IT-21	Information Technology for the 21st Century
ITIL	Information Technology Infrastructure Library
IWS	Interim Wideband System
kbps	1,000 bits per second
MAGTF	Marine Air-Ground Task Force

Mbps	1,000,000 bits per second
MCEITS	Marine Corps Enterprise IT Services
MCEN	Marine Corps Enterprise Network
MUOS	Mobile User Objective System
NECC	Network Enabled Command and Control
NetOps	Network Operations
NGEN	Next Generation Enterprise Network
NIAP	National Information Assurance Program
NIST	National Institute of Standards and Technology
NMCI	Navy Marine Corps Intranet
NNE	Naval Network Environment
NSF	National Science Foundation
OCONUS	Outside the Continental United States
ONE-NET	OCONUS Navy Enterprise Network
OPEX	Operating Expenditure
ORCA	Optical RF Communication Adjunct
PaaS	Platform as a Service
RF	Radio Frequency
SaaS	Software as a Service
SATCOM	Satellite Communications
SHF	Super High Frequency
SIPRnet	Secret Internet Protocol Router Network
SLA	Service Level Agreement
TOR	Terms of Reference
UHF	Ultra High Frequency
UHF F/O	UHF Follow-on (communications satellite system)
VL	Virtualization Layer
VM	Virtual Machines
WGS	Wideband Global SATCOM
XML	Extensible Markup Language